



Phishing for the Truth: A Scenario-Based Experiment of Users' Behavioural Response to Emails

Kathryn Parsons, Agata Mccormac, Malcolm Pattinson, Marcus Butavicius,
Cate Jerram

► To cite this version:

Kathryn Parsons, Agata Mccormac, Malcolm Pattinson, Marcus Butavicius, Cate Jerram. Phishing for the Truth: A Scenario-Based Experiment of Users' Behavioural Response to Emails. 28th Security and Privacy Protection in Information Processing Systems (SEC), Jul 2013, Auckland, New Zealand. pp.366-378, 10.1007/978-3-642-39218-4_27 . hal-01463838

HAL Id: hal-01463838

<https://inria.hal.science/hal-01463838>

Submitted on 9 Feb 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Phishing for the truth: A scenario-based experiment of users' behavioural response to emails

Kathryn Parsons ^a, Agata McCormac ^a, Malcolm Pattinson ^b, Marcus Butavicius ^a, Cate Jerram ^b

^a Defence Science and Technology Organisation

PO Box 1500, Edinburgh, SA, 5111, Australia

{Kathryn.Parsons; Agata.Mccormac; Marcus.Butavicius}@dsto.defence.gov.au

^b Business School, The University of Adelaide

10 Pulteney Street, Adelaide, SA, 5005, Australia

{Malcolm.Pattinson; Cate.Jerram}@adelaide.edu.au

· Corresponding author. Email: Kathryn.Parsons@dsto.defence.gov.au. Tel.: +61-8-7389-7953.
Fax: +61-8-7389-6328.

Phishing for the truth: A scenario-based experiment of users' behavioural response to emails

Using a role play scenario experiment, 117 participants were asked to manage 50 emails. To test whether the knowledge that participants are undertaking a phishing study impacts on their decisions, only half of the participants were informed that the study was assessing the ability to identify phishing emails. Results indicated that the participants who were informed that they were undertaking a phishing study were significantly better at correctly managing phishing emails and took longer to make decisions. This was not caused by a bias towards judging an email as a phishing attack, but instead, an increase in the ability to discriminate between phishing and real emails. Interestingly, participants who had formal training in information systems performed more poorly overall. Our results have implications for the interpretation of previous phishing studies, the design of future studies and for training and education campaigns, as it suggests that when people are primed about phishing risks, they adopt a more diligent screening approach to emails.

Keywords: phishing; information security; security behaviours; email security; security training

1. Introduction

Phishing is a term that describes an attempt to deceptively acquire personal and financial information via electronic communication with malicious intent. Social engineering strategies in conjunction with computer knowledge are used to gather

usernames, passwords, and bank account and credit card details (Anti-Phishing Working Group, 2010). Phishing attacks are commonly committed via email, and victims are often directed to fraudulent websites that appear legitimate (Moore & Clayton, 2007). Such breaches can have serious consequences, including direct consequences, such as financial loss if a phisher obtains access to a bank account, and indirect consequences, such as damaged reputation (Tam, Glassman & Vandenwauver, 2010).

Although there is a growing body of phishing studies, there is a lack of research examining the impact of the cognitive bias known as the subject expectancy effect (Anandpara, Dingman, Jakobsson, Liu & Roinestad, 2007). Essentially, studies where participants know they are participating in a phishing study have been criticised because they lack real world validity. This criticism is based on the assumption that individuals who are aware that they are taking part in a phishing study may be more suspicious, and this may therefore result in a bias towards 'phishing' decisions (Anandpara et al., 2007). It is unlikely that individuals would have this level of suspicion when checking their personal inboxes (Furnell, 2007).

In response to the possible influence of the subject expectancy effect, researchers have begun incorporating a role play scenario into the design of their phishing studies. This approach aims to minimise the bias caused by the subject expectancy effect, because participants are not informed that they are participating specifically in a phishing study. A study of this nature was conducted by Downs, Holbrook and Cranor (2007). Participants were informed that they were participating in a study about computer use, not computer security. They were given the identity of 'Pat Jones' and were shown images of emails from 'Pat's' inbox. Some emails were legitimate and some were phishing emails and participants were given options about how they would respond to each email. Downs et al. (2007) found that participants who were more knowledgeable and experienced with the internet environment were less susceptible to phishing attacks. They also found that participants' perceptions of the consequences of emails did not reliably predict their behaviour (2007). This study used only five email screen shots, which limits the ability to generalise these findings to other types of emails.

Downs, Holbrook and Cranor (2006) conducted a similar study that used a role play design and also incorporated qualitative interviews on computer security and trust. The study focused on decision strategies and susceptibility to phishing and concluded that participants were most likely to use subjective cues, such as relying on the text within an email, to determine the trustworthiness of emails, rather than relying on more objective cues, such as using information contained within URLs and links (Downs et al., 2006). They also found that participants were more vulnerable to unfamiliar phishing scams, and were generally less susceptible to scams they had seen previously. However, the authors acknowledge that their findings were significantly limited by a small sample size of only 20 participants (Downs et al., 2006), which, once again, limits the generalisability of findings.

The limitations and shortcomings of these previous studies provided the incentive for this current study, which was designed to address these issues, and further the vital work in the area of electronic mail fraud. For example, this is the first study that we are aware of that tested the influence of the subject expectancy effect on par-

ticipants' response to phishing emails. This was achieved by using a role play design where, although all of the participants were aware they were participating in a study about email, only half of the participants were informed they were participating in a *phishing* study. This will therefore reveal whether the knowledge that participants are participating in a phishing study influences their decisions. Our examination of the influence of the subject expectancy effect may highlight the need to reevaluate or interpret the findings of all previous phishing studies in light of the effect. This study will therefore provide vital knowledge regarding the design of future studies.

Participants in the current study were also exposed to a comparatively larger number of emails than in the previous studies, and these emails varied widely. Furthermore, to better understand what makes some people more susceptible to phishing emails than others, this study included a demographics questionnaire and a measure of impulsivity. This provided a more comprehensive assessment of individual differences than the previous research.

2. Methodology

2.1. Participants

A total of 117 students from the University of Adelaide were recruited via email and participated in the study. Of the 117 participants, 27 were male and 90 were female. The majority of participants were first year students (93), there were also 19 second year and 4 third year students, and 1 participant was completing post-graduate studies. Most participants were 25 years of age and younger (108) and only 5 participants were over the age of 30. Participants received \$25 cash for their participation.

2.2. Materials

2.2.1. Emails

The study consisted of 50 images of emails; half of these images were genuine emails, and half were phishing emails. The selected emails were comprised of 'actual' real and phishing emails that were either received by the authors, or found online. A range of emails, including banking, shopping and social networking emails were utilised to ensure that they were representative of the types of topics that would be expected in a typical inbox. An example of a phishing email can be seen in Figure 1, and an example of a real email used in the experiment can be seen in Figure 2.

A fictitious character, by the name of 'Sally Jones' was created, and the original emails that contained personal information were altered to include her details as if she was the intended recipient. Participants were informed that they were viewing emails from the inbox of 'Sally Jones', and were asked to make a decision regarding how they would manage each email. They were not provided with any other information regarding the persona of Sally Jones.



Fig. 1. Example of a phishing email. Email containing logo reprinted with permission from Google.



Fig. 2. Example of a real email. Email containing logo reprinted with permission from Google.

2.2.2. Measures of Individual Differences

A number of demographics were collected including information about gender and education level. The Cognitive Reflection Test (CRT), which is a very quick and efficient measure of impulsivity (Frederick, 2005), was also utilised. The test includes three questions and the most obvious response is not correct. To answer correctly, participants should stop and consider the question before providing an answer. A higher score on this test relates to a better ability to control impulsivity. This particular test was selected because findings indicate that the predictive validity of this measure was equal or above other cognitive measures (Frederick, 2005).

It is hypothesised that individuals who are better able to effectively manage impulsivity may be less susceptible to phishing emails, as they may be more likely to thoroughly deliberate the legitimacy of the email. This hypothesis was tested in a study by Kumaraguru and colleagues (2007). Participants with higher CRT scores were less likely to click on the phishing emails, but the results were not statistically significant. This may be due to the small number of phishing emails in the experiment, and the relatively small number of participants. The current study will retest this hypothesis with a larger number of emails and participants.

2.3. Method

Participants were informed that they were completing an experiment on how peo-

ple manage emails. They were told that they would be required to view images of 50 emails, taken from the inbox of Sally Jones. In order to test the influence of the subject expectancy effect, the participants were divided into two groups, the 'Control' Group and the 'Alerted' Group. The 'Control' Group, which consisted of 59 participants, were given the following description:

"Managers are often inundated with an extremely large number of emails on a daily basis, and the management of these emails is often very difficult. We're interested in assessing how people manage emails. You will be presented with a number of emails, both personal and work related, taken from the inbox of 'Sally Jones'.

Your job is to examine each email, with the aim of assisting Sally to process her Inbox. You will be asked what action you would recommend to her. You will also be asked to provide a rating of how confident you are with your recommendation, and what aspect of the email most influenced your recommendation."

The 'Alerted' Group, which consisted of 58 participants, were informed that they were participating in a phishing study and were given the same description with the following sentence added to the end:

"We are specifically interested in assessing the ability to identify 'phishing' emails. These are fraudulent email messages that are used to obtain personal information for the purposes of identity theft."

The research assistant also gave a verbal description of what phishing emails are to be sure that all participants had this knowledge.

For each of the 50 emails, all participants were asked to respond to the question, "How would you manage this email?" with one of four replies: a) leave the email in the inbox and flag for follow up; b) leave the email in the inbox; c) delete the email; or d) delete the email and block the sender. For each email, participants were also asked, "what aspect of this email influenced your decision?"

After responding to all 50 emails, participants were required to complete the demographics questionnaire and the cognitive test. The logic of this process ensured that the demographic questions could not alert the participants in the 'Control' Group that they were participating in a phishing study until after they had completed the main task.

3. Results

3.1. Phishing Emails

As shown in Figure 3, when responding to phishing emails, participants in the 'Control' Group most frequently responded with 'Flag for follow up' (37%), whereas the participants in the 'Alerted' Group were mostly likely to respond with 'Delete' (38%). The pie chart below also indicates that participants in the 'Alerted' Group were far more likely to respond with 'Delete and block' (11% of 'Control' Group responses versus 23% of 'Alerted' Group responses).




Fig. 1. Responses to phishing emails for the 'Control' and 'Alerted' Groups.

For the phishing emails, a response of 'Delete and block' was considered most appropriate, and a response of 'Flag for follow up' was deemed least appropriate. A total score was calculated for phishing emails, where a response of 'Delete and block' was assigned a score of 4, a response of 'Delete' was assigned a score of 3, a response of 'Leave in inbox' was assigned a score of 2, and a response of 'Flag for follow up' was assigned a score of 1. This assignment was such that the more appropriate the action when faced with a phishing email, the higher the values assigned to it.

A Mann-Whitney U-test was used to compare the ranks for the participants in the 'Control' Group and the 'Alerted' Group. The results indicated that there was a statistically significant difference, $U(116) = 2644.00$, $Z = 5.089$, $p < .001$. Participants in the 'Control' Group had a mean rank of 43.19, while participants in the 'Alerted'

Group had a mean rank of 75.09. This means that the participants in the ‘Alerted’ Group were significantly better at correctly managing the phishing emails, indicating that knowledge that participants were undertaking a phishing study tended to improve performance.

3.2. Real Emails

Interestingly, for the real emails, there was very little difference between the groups in regards to the frequency of both ‘Delete and block’ and ‘Flag for follow up’ responses. Instead, as shown in Figure 4, the percentage of responses in those categories was very similar, and the groups differed in regards to the responses, ‘Leave in inbox’ and ‘Delete’. When responding to real emails, the participants in the ‘Control’ Group were far more likely to respond with ‘Delete’ (37%), whereas the participants in the ‘Alerted’ Group were most likely to respond with ‘Leave in inbox’ (44%).




Fig. 2. Responses to real emails for the ‘Control’ and ‘Alerted’ Groups.

A total score was calculated for real emails using a similar approach to that employed above for the phishing emails, i.e., a response of ‘Flag for follow up’ was assigned a score of 4, a response of ‘Leave in inbox’ was assigned a score of 3, a response of ‘Delete’ was assigned a score of 2, and a response of ‘Delete and block’ was assigned a score of 1.

A Mann-Whitney *U*-test was used to compare the ranks for the participants in the ‘Control’ Group and the ‘Alerted’ Group. There were no statistically significant differences between the groups, $U(116) = 2008.50$, $Z = 1.624$, $p = .104$, with a mean rank for participants in the ‘Control’ Group of 53.96, and a mean rank for participants in the ‘Alerted’ Group of 64.13. Although the difference was not statistically significant, the participants who were alerted that they would be viewing some phishing emails were more likely to correctly manage the real emails.

3.3. Bias and Discrimination

To further examine the nature of any subject expectancy effect, the signal detection theory measures of discrimination and bias were calculated from the data (Green and Swets, 1966). In this context, discrimination refers to the ability of a participant to distinguish real from phishing emails while bias refers to an overall tendency to keep or delete emails in the inbox. A' and B'' were used as measures of discrimination and bias respectively (refer to Stanislaw and Todorov, 1999). These non-parametric measures are calculated directly from the commonly used measures 'Hit Rate' (HR) and 'False Alarm Rate' (FAR):

The HR refers to the probability that a phishing email is met with a phishing decision and the FAR refers to the probability that the participant responded with a phishing decision when it was a real email. For the purposes of this analysis, when faced with phishing emails, responses of 'Delete' or 'Delete and block' were deemed to be 'hits', and responses of 'Flag for follow up' or 'Leave in inbox' were deemed to be 'misses'. When faced with real emails, responses of 'Delete' or 'Delete and block' were deemed to be 'false alarms', and responses of 'Flag for follow up' or 'Leave in inbox' were deemed to be 'true misses'.

An A' value of 1 equates to perfect discrimination, and a value of 0.5 indicates that the respondent could not distinguish phishing emails from the real emails. A B'' value of zero indicates that there was no bias in the responses, a value of -1 indicates an extreme bias towards 'phishing' decisions, and a value of 1 indicates an extreme bias towards 'real' decisions.

The results indicated that the participants in the 'Alerted' Group were better able to discriminate between the phishing and real emails than the participants in the 'Control' Group ($A'_{\text{Control}} = 0.52$, $CI_{95\%} = [0.48, 0.56]$; $A'_{\text{Alerted}} = 0.72$, $CI_{95\%} = [0.68, 0.76]$). Furthermore, the results showed that participants in both the 'Control' and 'Alerted' Group had a very small response bias ($B''_{\text{Control}} = 0.07$, $CI_{95\%} = [0.01, 0.13]$; $B''_{\text{Alerted}} = 0.04$, $CI_{95\%} = [-0.04, 0.13]$), indicating that this study did not find evidence of the subject expectancy effect.

3.4. Time Taken

An independent-samples t-test demonstrated a significant difference between the time taken to manage emails between the two groups, $t(116) = 4.093$, $p < .001$. The participants in the 'Control' Group ($M = 21.47$, $SD = 5.83$) took significantly less time to make their decisions than the participants in the 'Alerted' Group ($M = 27.05$, $SD = 8.67$). The eta squared statistic (.13) indicated a large effect size. This therefore suggests that informing participants that they were completing a phishing study may have resulted in an increase in diligence and vigilance.

3.5. Individual Differences

Since the manner in which participants managed emails is best captured via the

four response categories, the mean ranks for phishing and real emails were used, and a series of Kruskal-Wallis and Mann-Whitney *U*-tests were conducted. The aim was to examine the influence of these variables on the mean ranks for phishing and real emails for the 'Control' Group and 'Alerted' Group.

3.5.1. Gender and Age

Contrary to the findings of Jagatic et al. (2005) and Sheng, Holbrook, Kumaraguru, Cranor and Downs (2010), who found that females and participants aged below 25 years were most vulnerable, the current study found no evidence of a relationship between either gender or age and the ability to correctly manage emails. However, of the 117 participants in this study, 90 were female and 108 were under 26 years of age. Because of this bias in our population, we can not discount the findings of Jagatic et al. (2005) and Sheng et al. (2010), and this is an issue worthy of further research.

3.5.2. Level of Education and Knowledge

After ranking the total phishing scores, a Kruskal-Wallis test was used to evaluate any association between performance in managing emails and participants' level of education. When participants were not told that they were conducting a phishing study, the participants with a higher level of education were significantly better at correctly managing phishing emails, $\chi^2 = 8.186$ (2, $N = 59$), $p = .017$ (Mean ranks; 'Year 12 or equivalent' = 27.33, 'Bachelor Degree' = 32.07, 'Honours Degree' = 56.17). However, there was no difference in the ability to correctly manage phishing emails when people were informed that they were completing a phishing study. This therefore suggests that people with more education were more likely to think about security without being prompted.

However, contrary to expectations, results indicated that the participants in the 'Control' Group who had completed a course in the area of information systems or information technology were less accurate in their ability to correctly manage phishing emails. A Mann-Whitney *U*-test revealed a mean rank of 21.50 for the $n = 17$ participants who had completed a course in the area of information systems or information technology, and a mean rank of 33.44 for the $n = 42$ who had not completed such a course, $U(58) = 501.50$, $Z = 2.421$, $p = .015$.

3.5.3. Employment Experience

A Kruskal-Wallis test revealed that, for the participants in the 'Alerted' Group, those who were currently employed (mean rank = 34.48) or had previous employment experience (mean rank = 30.69) were significantly better at identifying phishing emails than those without any employment experience (mean rank = 17.18), $\chi^2 = 7.817$ (2, $N = 58$), $p = .02$. Interestingly, this was only true for the participants who were informed that they were conducting a phishing study.

3.5.4. Cognitive Impulsivity

For participants in the 'Control' Group, those who obtained a higher score on the test of cognitive impulsivity (and were therefore better able to control impulsivity) were significantly better at identifying phishing emails, $\chi^2 = 8.241$ (3, $N = 59$), $p = .041$. The mean rank for the participants who obtained a score of three, which was the maximum possible score, was 48.92, which is significantly higher than the mean rank obtained by participants with the other possible scores (mean ranks; '0' = 27.89, '1' = 27.11, '2' = 29.79). This is consistent with our hypothesis. However, there were no significant differences for the participants in the 'Alerted' Group, which suggests that, once participants knew they were undertaking a phishing study, they were more likely to stop and consider their response, regardless of whether they were usually more impulsive.

3.6. Qualitative Content Analysis

Participants' responses to the question, "*What aspect of this email influenced your decision?*" were analysed using content analysis. An analysis of the participants' responses to the open-ended question in conjunction with their response to the multiple choice question supported the hypothesis that participants were more likely to 'Delete' or 'Delete and block' emails when they were suspicious of their legitimacy. In a minority of cases, participants flagged an email for follow up to alert 'Sally' of a possible security threat that she should report to the purported organisation. However, this reasoning was rare, and an examination of these cases revealed that re-categorising them would not impact on the statistical significance of the results reported above.

The justification of decisions supported the findings of Furnell, Tsaganidi and Phippen (2008), that participants were influenced by their perception of trust. This was based on the perceived trustworthiness of the company that the email appeared to originate from. For example, participants responded with statements such as "*[Company name] is a trusted chat program used all over the world, so emails from it would be legit*". These participants did not appear to question whether the email actually did originate from that company, but rather, appeared to decide based solely on the face validity of the email.

Many participants also mentioned the visual presentation of emails. For example, many participants deleted a real email from a telephone company because it did not contain any company logos, and therefore concluded that it seemed suspicious. This is consistent with the findings of Everard and Galletta (2006), who found the perceived quality of an online Web site was strongly influenced by the style of the Web site, and poor style was associated with low perceived quality.

Another common justification was based on incentives within emails. Participants were more susceptible to phishing emails when the email promised a financial reward. For example, two survey requests with a financial incentive for participation resulted in responses such as "*the \$100 monetary compensation is a great incentive for me to participate in this survey*". Hence, potential incentives may limit partici-

pants' ability to make valid and considered security decisions. It should be noted that participants were paid \$25 for their participation, and therefore, may have been more susceptible to offers of financial reward.

Other commonly cited reasons for decisions included spelling and grammatical errors, the personalisation of the email, and the perceived legitimacy of the URL. As also found by Furnell (2007), although these were cues that could conceivably prove useful, they often failed to assist participants in making the correct decisions concerning the legitimacy of an email. In support of previous research, which indicated that participants do not notice security indicators (Herzberg, 2009, Schechter et al., 2007), only one participant in our study used 'HTTPS' as a justification for their decision.

4. Discussion

This study provides further evidence that people are poor at identifying phishing emails. Overall, approximately 42% of all emails were incorrectly classified in this experiment. Participants who knew they were undertaking a phishing study were better able to make the distinction between real and phishing emails, which means that this study found no evidence of the subject expectancy effect. Participants in the 'Alerted' Group were not simply biased towards 'phishing' decisions, but were instead more likely to correctly manage all emails. Although the improved ability to correctly manage the real emails was not statistically significant, they were still more likely to correctly manage the real emails than the participants in the 'Control' Group. Evidence suggests that priming participants with the notion of phishing may have resulted in more diligent decision making, as the participants in the 'Alerted' Group took significantly longer to complete the experiment.

The influence of the different instructions provided to the two groups in our experiment may be explained by the general phenomenon known as framing (Tversky & Kahneman, 1981). In other words, it was the context in which the task was presented which influenced their decision making processes. Unlike the 'Control' group, the participants in the 'Alerted' group were specifically informed that the study was testing how well they could detect phishing emails and this had a positive influence on their decision-making. As discussed in the context of Signal Detection Theory (Green and Swets, 1966), this change was not in decision bias towards classifying an email as phishing but instead reflected an improvement in discrimination ability. Framing the task as one of detecting phishing emails may have caused participants to focus more on cues in the stimuli that better distinguished real from phishing emails. This has important implications for training and education programs, as it suggests that when people were primed to think about phishing, they were better able to identify phishing emails, and hence, less susceptible to phishing attacks.

The findings in regards to individual differences also have important implications for education and training programs. Participants who had attended an information systems or technology course were, in fact, less likely to correctly manage emails. This may suggest that knowledge in this area could lead to complacency.

Instead, actual security behaviours (such as using spam filters and adjusting security preferences) were better predictors of the ability to deal with phishing emails. In addition, when the task was framed as a phishing test, participants with employment experience performed better, possibly as a result of more experience in dealing with categorising emails in a work environment. Our results also indicated that participants who were better able to control impulsivity were better at managing phishing emails. This suggests that it may be more effective to emphasise the importance of stopping and thinking before responding to any email rather than exclusively teaching security rules. This is supported by the training literature, which indicates that it is more effective to emphasise specific behaviours rather than rules (Parsons, McCormac, Butavicius & Ferguson, 2010).

Our findings also have implications for the research literature on users' susceptibility to phishing emails. Previous studies should be interpreted in the light of the 'framing' effect identified in this study and future research should carefully consider how the task is presented to the user. Critically, the discovery of the framing effect suggests that the risk of phishing may, on the whole, be underestimated in previous literature. Specifically, in our study the inferior results of the 'Control' Group, who were not informed that they were undertaking a phishing study, better represent the performance of real-world users. This is because in real life the frequency with which people are reminded about the risks of phishing emails is generally low.

It is, however, important to highlight the fact that phishing studies such as ours do not directly measure actual susceptibility. In our experiment, participants were not required to click on any of the links or provide personal information, and it is therefore possible that, in a real world situation, participants may have become suspicious before succumbing to any of the phishing attacks. This study was also a role play, and the manner in which participants deal with emails in an experimental environment may not relate precisely to how participants would deal with actual emails received in their personal inboxes. Furthermore, in this study, participants did not know which sites 'Sally' subscribes to, and therefore their ability to make context dependent decisions was limited. In a real life situation, whether someone is a member of a particular bank or social networking site is likely to influence the decision to delete or keep an email. Future research should investigate how to more accurately replicate these variables in an experimental context.

References

- Anandpara V, Dingman A, Jakobsson M, Liu D, Roinestad H. Phishing IQ tests measure fear, not ability. In: Proceedings of the 11th International Conference on Financial cryptography and 1st International conference on Usable Security, Scarborough, Trinidad and Tobago; 2007. p. 362–66.
- Anti-Phishing Working Group. Global Phishing Survey: Trends and Domain Name Use in 2H2009. Available from: <<http://www.antiphishing.org/>>; May 2010.
- Downs JS, Holbrook MB, Cranor LF. Decision strategies and susceptibility to phishing. In: Proceedings of the Second Symposium On Usable Privacy and Security, Pittsburgh, PA, USA; 2006. p. 79–90.

- Downs JS, Holbrook M, Cranor LF. Behavioral response to phishing risk. In: Proceedings of the Anti-Phishing Working Groups 2nd Annual eCrime Researchers Summit, Pittsburgh, PA, USA; 2007. p. 37–44.
- Frederick S. Cognitive reflection and decision making. *Journal of Economic Perspectives*. 2005;16(4):25–42.
- Furnell S. Phishing: can we spot the signs? *Computer Fraud & Security*. 2007;3:10–15.
- Furnell S, Tsaganidi V, Phippen A. Security beliefs and barriers for novice Internet users. *Computers & Security*. 2008;27:235–40.
- Green DM, Swets JA. *Signal Detection Theory and Psychophysics*. New York: Wiley; 1966.
- Herzberg A. Why Johnny can't surf (safely)? Attacks and defenses for web users. *Computers & Security*. 2009;28:63–71.
- Jagatic TN, Johnson NA, Jakobsson M, Menczer F. Social phishing. *Communications of the ACM*. 2007;50(10):94–100.
- John OP, Donahue EM, Kentle RL. *The Big Five Inventory--Versions 4a and 54*. Berkeley, CA: University of California, Berkeley, Institute of Personality and Social Research; 1991.
- John OP, Naumann LP, Soto CJ. Paradigm shift to the integrative big-five trait taxonomy: History, measurement, and conceptual issues. In John OP, Robins RW, Pervin, LA, editors. *Handbook of personality: Theory and research* (3rd ed.). New York, NY: Guilford Press; 2008. p. 114–58.
- John OP, Srivastava S. The big-five trait taxonomy: History, measurement, and theoretical perspectives. In Pervin LA, John OP, editors. *Handbook of personality: Theory and research* (2nd ed.). New York: Guilford Press; 1999. p. 102–39.
- Kumaraguru P, Rhee Y, Sheng S, Hasan S, Acquisti A, Cranor LF, Hong J. Getting users to pay attention to anti-phishing education: Evaluation of retention and transfer. In: Proceedings of the 2nd Annual eCrime Researchers Summit, Pittsburgh, PA; 2007. p. 70–81.
- Moore T, Clayton R. An empirical analysis of the current state of phishing attack and defence. In: Proceedings of the Sixth Workshop on the Economics of Information Security, Pittsburgh, PA, USA; 2007. p. 1–20.
- Parsons K, McCormac A, Butavicius M, Ferguson L. *Human Factors and Information Security: Individual, Culture and Security Environment*. DSTO Technical Report, DSTO-TR2484; 2010.
- Sheng S, Holbrook M, Kumaraguru P, Cranor L, Downs J. Who falls for phish? A demographic analysis of phishing susceptibility and effectiveness of interventions. In: Proceedings of the 28th International Conference on Human Factors in Computing Systems, Atlanta, Georgia, USA; 2010. p. 373–82.
- Stanislaw H, Todorov N. Calculation of signal detection theory measures. *Behavior Research Methods Instruments & Computers*. 1999;31(1):137–49.
- Tam L, Glassman M, Vandenwauver M. The psychology of password management: a tradeoff between security and convenience. *Behaviour & Information Technology*. 2010;29(3):233–44.
- Tversky A, Kahneman D. The framing of decisions and the psychology of choice, *Science*. 1981;185:453–58.